

Distributed Randomized Control for Demand Dispatch

Ana Bušić and Sean Meyn[†]

March 21, 2016

Abstract

The paper concerns design of control systems for *Demand Dispatch* to obtain ancillary services to the power grid by harnessing inherent flexibility in many loads. The role of “local intelligence” at the load has been advocated in prior work; randomized local controllers that manifest this intelligence are convenient for loads with a finite number of states. The present work introduces two new design techniques for these randomized controllers:

- (i) The *Individual Perspective Design* (IPD) is based on the solution to a one-dimensional family of Markov Decision Processes, whose objective function is formulated from the point of view of a single load. The family of dynamic programming equation appears complex, but it is shown that it is obtained through the solution of a single ordinary differential equation.
- (ii) The *System Perspective Design* (SPD) is motivated by a single objective of the grid operator: Passivity of any linearization of the aggregate input-output model. A solution is obtained that can again be computed through the solution of a single ordinary differential equation.

Numerical results complement these theoretical results.

1 Introduction

Renewable energy sources such as wind and solar power have a high degree of unpredictability and time variation, which makes balancing demand and supply increasingly challenging. One way to address this challenge is to harness the inherent flexibility in demand of many types of loads. Loads can supply a range of ancillary services to the grid, such as the balancing reserves required at Bonneville Power Authority (BPA), or the Reg-D/A regulation reserves used at PJM [1]. Today these services are secured by a balancing authority (BA) in each region.

1.1 Demand dispatch

These grid services can be obtained without impacting quality of service (QoS) for consumers [1, 2], but this is only possible through design. The term *Demand Dispatch* is used in this paper to emphasize the difference between the goals of our own work and traditional demand response.

Consumers use power for a reason, and expect some guarantees on the QoS they receive. The grid operator desires reliable ancillary service, obtained from the inherent flexibility in the consumer’s power consumption. These seemingly conflicting goals can be achieved simultaneously, but

*Research supported by French National Research Agency grant ANR-12-MONU-0019, and NSF grants CPS-0931416 and CPS-1259040.

[†]A.B. is with Inria and the Computer Science Dept. of École Normale Supérieure, Paris, France; S.M. is with the Department of Electrical and Computer Engg. at the University of Florida, Gainesville.




Figure 1: Wind generation at BPA.

the solution requires local control: an appliance must monitor its QoS and other state variables, it must receive grid-level information (e.g., from the BA), and based on this information it must adjust its power consumption. With proper design, an aggregate of loads can be viewed at the grid-level as *virtual energy storage* (VES). Just like a battery, the aggregate provides ancillary service, even though it cannot produce energy.

Fig. 1 shows the wind generation in the BPA region during the first week of 2015. There is virtually no power generated on New Year’s Day, and generation ramps up to nearly 4GW on the morning of January 5. In this example we show how to supply a demand of exactly 4GW during this time period, using generation from wind and other resources.




Figure 2: One power signal representing significant energy delivery, and two zero-energy power signals combine with wind generation to match a 4GW load.

Let $G_r(t)$ denote the additional power required at time t , in units of GWs. For example, on the first day of this week we have $G_r(t) \approx 4$. Day-ahead forecast of the low frequency component of generation from wind is highly predictable. We let G_{LP} denote the signal obtained by passing the forecast of G_r through a low pass filter. The signal G_{HP} is obtained by filtering $G_r - G_{LP}$ using a high pass filter, and $G_{MP} = G_r - G_{LP} - G_{HP}$. Each of these filters is causal, and their parameters are a design choice. Examples are shown in Fig. 2.

It is not difficult to ramp hydro-generation up and down to accurately track the power signal G_{LP} . This is an energy product that might be secured in today’s day-ahead markets.

The other two signals shown in Fig. 2 take on positive and negative values. Each represents a total energy of approximately zero, hence it would be a mistake to attempt to obtain these services in an energy market. Either could be obtained from a large fleet of batteries or flywheels. However, it may be much cheaper to employ flexible loads via demand dispatch. The signal G_{HP} can be obtained by modulating the fans in commercial buildings (perhaps by less than 10%) [3]. The signal G_{MP} can be supplied in whole or in part by loads such as water heaters, commercial refrigeration, and water chillers.

Low frequency variability from solar gives rise to the famous “duck curve” anticipated at

CAISO¹, which is represented as the hypothetical “net-load curve” in Fig. 3a. The actual net-load curve is the difference between load and generation from renewables; the drop from 20 to 10 GW is expected with the introduction of 10 GW of solar in the state of California. This curve highlights the ramping pressure placed on conventional generation resources.

As shown in this figure, the volatility and steep ramps associated with California’s duck curve can be addressed using a frequency decomposition: The plot shows how the net-load can be expressed as the sum of four signals distinguished by frequency. Variability introduced by the low frequency component can be rejected using traditional resources such as thermal generators, along with some flexible loads (e.g., from flexible industrial manufacturing). The mid-pass signal shown in the figure would be a challenge to generators for various reasons, but this zero-mean signal, as well as the higher frequency components, can be tracked using a variety of flexible loads.

The control architecture described in this paper is not limited to handling disturbances from wind and solar energy. Fig. 3b illustrates how the same frequency decomposition can be used to allocate resources following an unexpected contingency, such as a generator outage.

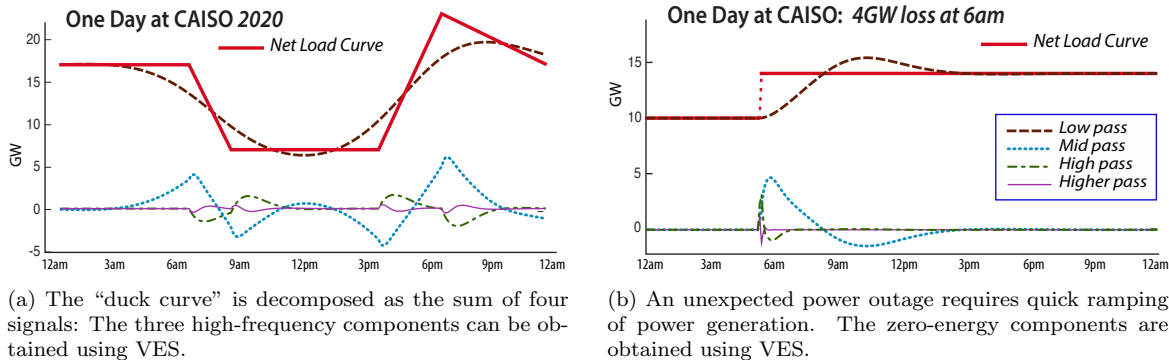


Figure 3: Frequency decomposition in two examples in which large ramps in generation are required. In each case, the low-pass energy component can be obtained through conventional generation, or load-shedding. The remaining components of net-load are obtained using VES.

For loads whose power consumption cannot be varied continuously, we have argued in prior work that a distributed randomized control architecture is convenient for design [2, 4, 5]. This architecture includes local control to maintain bounds on the quality of service delivered by the loads, and also to ensure high quality ancillary service to the grid.

Analysis of the aggregate is based on a mean-field model.

1.2 Mean-field model

We restrict to the setting of the prior work [2, 4], based on the following distributed control architecture. A family of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$ is constructed to define local decision making. Each load evolves as a controlled Markov chain on a finite state space, with common input $\zeta = (\zeta_0, \zeta_1, \dots)$. It is assumed that the scalar signal ζ is broadcast to each load. If a load is in state x at time t , and the value ζ_t is broadcast, then the load transitions to the state x' with probability $P_{\zeta_t}(x, x')$. Letting X_t^i denote the state of the i th load at time t , and assuming N loads, the empirical pdf (probability mass function) is defined as the average,

$$\mu_t^N(x) = \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{X_t^i = x\}, \quad x \in \mathbb{X}.$$

¹an ISO in California: www.caiso.com

The mean-field model is the deterministic system defined by the evolution equations,

$$\mu_{t+1} = \mu_t P_{\zeta_t}, \quad t \geq 0, \quad (1)$$

in which μ_t is a row vector. Under general conditions on the model and on μ_0 it can be shown that μ_t^N is approximated by μ_t .

In this prior work it is assumed that average power consumption is obtained through measurements or state estimation: Assume that $\mathcal{U}(x)$ is the power consumption when the load is in state x , for some function $\mathcal{U}: \mathbf{X} \rightarrow \mathbb{R}_+$. The average power consumption is defined by,

$$y_t^N = \frac{1}{N} \sum_{i=1}^N \mathcal{U}(X_t^i), \quad x \in \mathbf{X}.$$

which is approximated using the mean-field model:

$$y_t = \sum_x \mu_t(x) \mathcal{U}(x), \quad t \geq 0. \quad (2)$$

The mean-field model is a state space model that is linear in the state μ_t , and nonlinear in the input ζ_t . The observation process (2) is also linear as a function of the state. Assumptions imposed in the prior work [2, 4, 5] imply that the input is a continuous function of these values.

In [4], the design of the feedback law $\zeta_t = \phi_t(y_0, \dots, y_t)$ is based on a linearization of this state space model. One goal of the present paper is to develop design techniques to ensure that the linearized input-output model has desirable properties for control design at the grid level.

1.3 Contributions

Several new design techniques are introduced in this paper, and the applications go far beyond prior work:

- (i) *Optimal design.* In the prior work [4], the family of transition matrices $\{P_\zeta\}$ was constructed based on an average-cost optimal control problem (MDP). The cost function in this MDP was parameterized by the scalar ζ . In this prior work, the optimal control problem was completely unconstrained, in the sense that any choice of P was permissible in the optimal control formulation. The optimal control formulation proposed here is far more general: We allow some randomness by design, and some exogenous randomness that is beyond our control. These contributions are summarized in Theorem 2.2.
- (ii) *Passivity by design.* A discrete-time transfer function F is *positive-real* if it is stable (all poles are strictly within the unit disk), and the following bound holds:

$$F(e^{j\theta}) + F(e^{-j\theta}) \geq 0, \quad \theta \in \mathbb{R}$$

It is *strictly positive-real* if the inequality is strict for all θ . A linear system is passive if it is positive real.

Let (A, B, C) describe a state-space representative of the linearization, with transfer function $G(z) = C(Iz - A)^{-1}B$. Consider the delay-free model with transfer function $G^+(z) := zG(z)$. That is,

$$G^+(z) = zC[Iz - A]^{-1}B = \sum_{k=0}^{\infty} CA^k Bz^{-k} \quad (3)$$

A new approach to design of the transition matrices is introduced in this paper to ensure that the linearization is strictly positive real. The main conclusions are summarized in Theorem 2.3.

- (iii) *ODE methods for design.* A unified computational framework is introduced. The construction of the transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$ is obtained as the solution to a single ODE for each of the design techniques (i) and (ii).
- (iv) *Applications.* Prior work on distributed control for demand dispatch focused on a single collection of loads: residential pool pumps [4–6]. The motivation was ease of exposition, and also the fact that the design methodology required special assumptions on nominal behavior. The new methodology developed in this paper relaxes these assumptions, and allows application to any load with discrete power states, such as a refrigerator, or other thermostatically controlled loads (TCLs).

1.4 Prior work

There are many recent papers with similar goals – to create a science to support demand dispatch. In [7] and its sequels, all control decisions are at the balancing authority, and this architecture then requires state estimation to obtain the grid-level control law. A centralized deterministic approach is developed in [8, 9]. None of this prior work considers design of local control algorithms, which is the focus of this paper.

Passivity was established in the prior work [10], but only for continuous time models for which the nominal model (with $\zeta = 0$) is a *reversible* Markov process. It follows that G^+ is minimum phase, and hence the original transfer function $G(z) = z^{-1}G^+(z)$ is also minimum phase.

The remainder of the paper is organized into four sections. Several design techniques are introduced in Section 2, and Section 3 presents general methodology for constructing and analyzing the linearized mean-field dynamics. Examples are contained in Section 4, and conclusions in Section 5.

2 Design

We first summarize the assumptions and notation.

2.1 Assumptions and Notation

A Markovian model for an individual load is created based on its typical operating behavior. This is modeled by a Markov chain with transition matrix denoted P_0 , with state space $\mathbf{X} = \{x^1, \dots, x^d\}$; it is assumed to be *irreducible and aperiodic*. It follows that P_0 admits a unique invariant pmf (probability mass function), denoted π_0 , and satisfying $\pi_0(x) > 0$ for each x .

It is assumed throughout this paper that the family of transition matrices used for distributed control is of the form,

$$P_\zeta(x, x') := P_0(x, x') \exp(h_\zeta(x, x') - \Lambda_{h_\zeta}(x)) \quad (4)$$

in which h_ζ is continuously differentiable in ζ , and Λ_{h_ζ} is the normalizing constant

$$\Lambda_{h_\zeta}(x) := \log\left(\sum_{x'} P_0(x, x') \exp(h_\zeta(x, x'))\right) \quad (5)$$

Each P_ζ must also be irreducible and aperiodic.

For any transition matrix P , an invariant pmf is interpreted as a row vector, so that invariance can be expressed $\pi P = \pi$. Any function $f: \mathbf{X} \rightarrow \mathbb{R}$ is interpreted as a d -dimensional column vector, and we use the standard notation $Pf(x) = \sum_{x'} P(x, x')f(x')$, $x \in \mathbf{X}$.

Several other matrices are defined based on P and π : The *adjoint* of P (in $L_2(\pi)$) is the transition matrix defined by

$$P^\natural(x, x') = \frac{\pi(x')}{\pi(x)} P(x', x), \quad x, x' \in \mathbf{X}. \quad (6)$$

The *fundamental matrix* is the inverse

$$Z = [I - P + 1 \otimes \pi]^{-1} = \sum_{n=0}^{\infty} [P - 1 \otimes \pi]^n \quad (7)$$

with $[P - 1 \otimes \pi]^0 := I$ (the $d \times d$ identity matrix), $1 \otimes \pi$ is a matrix in which each row is identical, and equal to π , and $[P - 1 \otimes \pi]^n = P^n - 1 \otimes \pi$ for $n \geq 1$.

The *Donsker-Varadhan rate function* is denoted,

$$K(P \| P_0) = \sum_{x, x'} \pi(x) P(x, x') \log \left(\frac{P(x, x')}{P_0(x, x')} \right) \quad (8)$$

It is used here to model the cost of deviation from the nominal transition matrix P_0 , as in [4, 10–12].

Nature & nurture In many applications it is necessary to include a model of randomness from nature along with the randomness introduced by the local control algorithm (nurture).

Consider a load model in which the full state space is the cartesian product of two finite state spaces: $\mathbf{X} = \mathbf{X}_u \times \mathbf{X}_n$, where \mathbf{X}_u are components of the state that can be directly manipulated through control. The “nature” components \mathbf{X}_n are not subject to direct control. For example, these components may be used to model the impact of the weather on the climate of a building.

Elements of \mathbf{X} are denoted $x = (x_u, x_n)$. Any state transition matrix under consideration is assumed to have the following conditional-independence structure,

$$P(x, x') = R(x, x'_u) Q_0(x, x'_n), \quad (9)$$

for $x \in \mathbf{X}$, $x'_u \in \mathbf{X}_u$, $x'_n \in \mathbf{X}_n$, where $\sum_{x'_u} R(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) = 1$ for each x . The matrix Q_0 is out of our control – this models load dynamics and exogenous disturbances.

2.2 Common structure for design

The construction of the family of functions $\{h_\zeta : \zeta \in \mathbb{R}\}$ in (4) is achieved using the following steps.

Step 1: The specification of a function \mathcal{H} that takes as input a transition matrix P that is irreducible. The output $H = \mathcal{H}(P)$ is a real-valued function on the product space $\mathbf{X} \times \mathbf{X}$. That is, $H(x, x') \in \mathbb{R}$ for each pair $(x, x') \in \mathbf{X} \times \mathbf{X}$.

Step 2: The family of transition matrices $\{P_\zeta\}$ and functions $\{h_\zeta\}$ are defined by the solution to the d -dimensional ODE:

$$\frac{d}{d\zeta} h_\zeta = \mathcal{H}(P_\zeta), \quad \zeta \in \mathbb{R}, \quad (10)$$

in which P_ζ is determined by h_ζ through (4). The boundary condition for this ODE is $h_0 \equiv 0$.

In the special case in which randomness from nature is not considered, we can apply the methods described here using $\mathbf{X}_n = \{x_n^1\}$ (a singleton).

The conditional independence constraint (9) imposes constraints on the functions $\{h_\zeta\}$ and the transformation \mathcal{H} . To ensure that P_ζ is of the form (9), it is sufficient to restrict to functions h_ζ of

(x, x') that do not depend on x'_n , where $x' = (x'_u, x'_n) \in \mathbf{X}$. For this reason we make the notational convention,

$$h_\zeta(x, x') = h_\zeta(x, x'_u), \quad x \in \mathbf{X}, \quad x' = (x'_u, x'_n) \in \mathbf{X}.$$

Since h_ζ is constructed through the ODE (10), we impose the same constraints on $H = \mathcal{H}(P)$:

$$H(x, x') = H(x, x'_u), \quad x \in \mathbf{X}, \quad x' = (x'_u, x'_n) \in \mathbf{X}.$$

Given any function $H^\circ: \mathbf{X} \rightarrow \mathbb{R}$, the function defined below satisfies this constraint:

$$H(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) H^\circ(x'_u, x'_n) \quad (11)$$

Each of the methods that follow construct $H = \mathcal{H}(P)$ of this form. Hence the design problem reduces to choosing a mapping $H^\circ = \mathcal{H}^\circ(P)$.

The following normalization is imposed throughout: The transition matrix P_ζ defined in (4) does not change if we add a constant to the function h_ζ . We are thus free to normalize $H^\circ = \mathcal{H}^\circ(P)$ by a constant. Throughout the paper we fix a state $x^\circ \in \mathbf{X}$, and design \mathcal{H}° so that $H^\circ(x^\circ) = 0$ for any P .

The ODE method can be simplified based on these observations. The proof of Prop. 2.1 is straightforward.

Proposition 2.1. *Consider a solution to the ODE (10) in which $H = \mathcal{H}(P)$ is of the form (11) for any matrix P . It then follows that each of the functions $\{h_\zeta\}$ are of this form:*

$$h_\zeta(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) h_\zeta^\circ(x'_u, x'_n),$$

$x \in \mathbf{X}$, $x' = (x'_u, x'_n) \in \mathbf{X}$, for some $h_\zeta^\circ: \mathbf{X} \rightarrow \mathbb{R}$. Moreover, these functions solve the d -dimensional ODE,

$$\frac{d}{d\zeta} h_\zeta^\circ = \mathcal{H}^\circ(P_\zeta), \quad \zeta \in \mathbb{R},$$

in which P_ζ is determined by h_ζ through (4), and with boundary condition $h_0^\circ \equiv 0$. ■

2.3 Individual Perspective

In this design, the mapping $H^\circ = \mathcal{H}^\circ(P)$ is defined in terms of the fundamental matrix:

IPD solution: Given P , the fundamental matrix Z is obtained from (7), and then for each $x \in \mathbf{X}$,

$$H^\circ(x) = \sum_{x'} [Z(x, x') - Z(x^\circ, x')] \mathcal{U}(x'). \quad (12)$$

The function H° specified in (12) is a solution to Poisson's equation,

$$PH^\circ = H^\circ - \mathcal{U} + \bar{\mathcal{U}} \quad (13)$$

where $\bar{\mathcal{U}}$ (also written $\pi(\mathcal{U})$) is the steady-state mean:

$$\bar{\mathcal{U}} := \sum_x \pi(x) \mathcal{U}(x) \quad (14)$$

The function (12) is the unique solution satisfying $H^\circ(x^\circ) = 0$ [13, Thm. 17.7.2].

This choice for H° is called the Individual Perspective Design (IPD) since h_ζ solves an optimization problem formulated from the point of view of a single load. Given $\zeta \in \mathbb{R}$, the “optimal reward” is defined by the maximum,

$$\eta_\zeta^* = \max_{\pi, P} \{ \zeta \pi(\mathcal{U}) - K(P \| P_0) : \pi P = \pi \} \quad (15)$$

and P is also subject to the structural constraint (9). The maximizer defines a transition matrix that is denoted,

$$\check{P}_\zeta = \arg \max_P \{ \zeta \pi(\mathcal{U}) - K(P \| P_0) : \pi P = \pi \} \quad (16)$$

It is shown in Theorem 2.2 that the optimal value η_ζ^* together with a *relative value function* h_ζ^* solve the average reward optimization equation (AROE):

$$\max_P \left\{ \mathcal{W}_\zeta(x, P) + \sum_{x'} P(x, x') h_\zeta^*(x') \right\} = h_\zeta^*(x) + \eta_\zeta^* \quad (17)$$

where $\mathcal{W}_\zeta(x, P) = \zeta \mathcal{U}(x) - \sum_{x'} P(x, x') \log \left(\frac{P(x, x')}{P_0(x, x')} \right)$.

The relative value function is not unique, since we can add a constant to obtain a new solution. We normalize this function so that $h_\zeta^*(x^\circ) = 0$. The proof of the following can be found in the working paper [14].

Theorem 2.2. *The IPD solution results in a collection of transition matrices $\{P_\zeta : \zeta \in \mathbb{R}\}$ with the following properties. For each ζ ,*

- (i) *The transition matrix is optimal, $P_\zeta = \check{P}_\zeta$.*
- (ii) *For each x and x'_u , the function h_ζ that defines P_ζ is given by,*

$$h_\zeta(x, x'_u) = \sum_{x'_n} Q_0(x, x'_n) h_\zeta^*(x'_u, x'_n) \quad (18)$$

where $(h_\zeta^, \eta_\zeta^*)$ solves the AROE (17).*

- (iii) *The steady-state mean power consumption satisfies,*

$$\frac{d}{d\zeta} \bar{u}_\zeta = \frac{d^2}{d\zeta^2} \eta_\zeta^* \geq 0 \quad (19)$$

and hence \bar{u}_ζ is monotone in ζ . ■

2.4 System Perspective

The motivation for the following System Perspective Design (SPD) is from the point of view of the BA. Under general conditions, the linearized aggregate model is passive, which is a desirable property from the grid-level perspective.

The construction of $H^\circ = \mathcal{H}^\circ(P)$ is similar to IPD. For any matrix P with invariant pmf π , recall the definition of the adjoint P^∇ in (6). The matrix product is denoted

$$P^\nabla(x, x') = \sum_{z \in \mathbf{X}} P^\nabla(x, z) P(z, x'), \quad x, x' \in \mathbf{X}.$$

The fundamental matrix defined in terms of this transition matrix is denoted $Z^\nabla = [I - P^\nabla + \mathbf{1} \otimes \pi]^{-1}$.

SPD solution: Given P , the matrix Z^∇ is obtained, and

$$H^\circ(x) = \sum_{x'} [Z^\nabla(x, x') - Z^\nabla(x^\circ, x')] \mathcal{U}(x'), \quad x \in \mathsf{X}. \quad (20)$$

Under additional assumptions, the algorithm obtained from SPD results in a positive real linearization. The proof of the bound (21) is contained in Section 3.2.

Theorem 2.3. *Suppose that the Markov chain with transition matrix $P_0^\nabla = P_0^\dagger P_0$ is irreducible, and that $P_0 = R_0$ (a model without probabilistic constraints).*

Then, the solution to the SPD satisfies the following strict positive-real condition: the linearized model at any constant value ζ obeys the bound,

$$G_\zeta^+(e^{j\theta}) + G_\zeta^+(e^{-j\theta}) \geq \sigma_\zeta^2, \quad \theta \in \mathbb{R} \quad (21)$$

where σ_ζ^2 is the variance of \mathcal{U} under π_ζ . ■

The irreducibility assumption on P_0^∇ does not come for free. Consider for example the Markov chain on $d \geq 3$ states defined by $P_0(x^i, x^{i+1}) = 1$ for $1 \leq i \leq d-1$, and $P_0(x^d, x^d) = P_0(x^d, x^1) = 1/2$. This chain is irreducible and aperiodic. The behavior of the adjoint is similar; in particular, $P_0^\dagger(x^{i+1}, x^i) = 1$ for each $1 \leq i \leq d-1$. It follows that $P_0^\nabla(x^k, x^k) = 1$ for $2 \leq k \leq d$, so the irreducibility assumption fails.

2.5 Exponential family

Rather than solve an ODE, it is natural to fix a function $H_e^\circ: \mathsf{X} \rightarrow \mathbb{R}$, and define for each x, x'_u and ζ ,

$$h_\zeta(x, x'_u) = \zeta H_e(x'_u | x)$$

$$\text{with } H_e(x'_u | x) := \sum_{x'_n} Q_0(x, x'_n) H_e^\circ(x'_u, x'_n)$$

This is a special case of the two-step design described in Section 2.2 in which $\mathcal{H}^\circ(P) = H_e^\circ$, independent of P , and the function $H = \mathcal{H}(P)$ is then obtained from (11).

In this case, the transition matrices defined in (4) can be regarded as an *exponential family*. The exponential family using $H_e^\circ = \mathcal{U}$ will be called the *myopic design*.

Other designs can be obtained as linear approximations to the IPD or SPD solutions, with $H_e^\circ = \mathcal{H}^\circ(P_0)$. In the linear approximation of the IPD solution, this is a solution Poisson's equation for the nominal model:

$$P_0 H_e^\circ = H_e^\circ - \tilde{\mathcal{U}}_0 \quad (22)$$

where $\tilde{\mathcal{U}}_0(x) = \mathcal{U}(x) - \bar{\mathcal{U}}_0$. The resulting exponential family is called the IPD_0 design. It is approximately optimal for ζ near zero – a proof of Theorem 2.4 can be found in [14].

Theorem 2.4. *The following approximations hold for the transition matrices $\{P_\zeta\}$ obtained from the IPD_0 design:*

(i) *With \check{P}_ζ the optimal transition matrix in (16),*

$$P_\zeta(x, x') = \check{P}_\zeta(x, x') + O(\zeta^2), \quad \text{for all } x, x', \zeta \quad (23)$$

(ii) Let $\eta_\zeta = \zeta \pi_\zeta(\mathcal{U}) - K(P_\zeta \| P_0)$ denote the value of the quantity in brackets in (15) that is obtained using (π_ζ, P_ζ) . Then, $\eta_\zeta = \eta_\zeta^* + O(\zeta^4)$. ■

A similar result holds if H_e° is chosen based on SPD, with

$$H_e^\circ = [I - P_0^\nabla + 1 \otimes \pi_0]^{-1}$$

The linearization at $\zeta = 0$ will be positive-real under the assumptions of Theorem 2.3, because (21) continues to hold at $\zeta = 0$,

$$G_0^+(e^{j\theta}) + G_0^+(e^{-j\theta}) \geq \sigma_0^2, \quad \theta \in \mathbb{R}.$$

An example in Section 4.2 shows that passivity may fail for the linearization G_ζ^+ at values of ζ far from zero.

2.6 Geometric sampling

Geometric sampling is specified by a transition matrix S_0 and a fixed parameter $\gamma \in (0, 1)$. At each time t , a weighted coin is flipped with probability of heads equal to γ . If the outcome is a tail, then the state does not change. Otherwise, a transition is made from the current state x to a new state x' with probability $S_0(x, x')$. The overall transition matrix is expressed as a convex combination,

$$P_0 = (1 - \gamma)I + \gamma S_0 \tag{24}$$

One motivation for sampling in [5] is to reduce the chance of excessive cycling at the loads, while ensuring that the data rate from balancing authority to loads is not limited. It was also found that this architecture justified a smaller state space for the Markov model.

Based on this nominal model, there are two approaches to applying the design techniques introduced in this paper. If P_0 is transformed directly, then the resulting family of transition matrix will be of the form,

$$P_\zeta = (1 - \gamma_\zeta)I + \gamma_\zeta S_\zeta \tag{25}$$

in which γ_ζ is a function of x . That is, if at time t the state is $X(t) = x$ and the input $\zeta_t = \zeta$, then once again a weighted coin is flipped, but with probability of success equal to $\gamma_\zeta(x)$. Conditioned on success, a transition is made to state x' with probability $S_\zeta(x, x')$.

In some cases it is convenient to fix the statistics of the sampling process, and transform S_0 using any of the design techniques described in the previous subsections. Once the family of transition matrices $\{S_\zeta : \zeta \in \mathbb{R}\}$ is constructed, we then define

$$P_\zeta = (1 - \gamma)I + \gamma S_\zeta \tag{26}$$

Each approach is illustrated through examples in Section 4.

3 Linearized Mean-Field Model

In this section we describe structure for the linearized model in full generality. We consider a general family of transition matrices of the form (4), maintaining the assumption that h_ζ is continuously differentiable in ζ , and that P_0 is irreducible and aperiodic.

3.1 Transfer function

Representations of the transfer function for the linearization require a bit more notation. We denote $\tilde{\mathcal{U}}_\zeta = \mathcal{U} - \bar{\mathcal{U}}_\zeta$, with $\bar{\mathcal{U}}_\zeta = \pi_\zeta(\mathcal{U})$. The derivative of the transition matrix is also a $d \times d$ matrix, denoted

$$\mathcal{E}_\zeta = \frac{d}{d\zeta} P_\zeta \quad (27)$$

A simple representation for this matrix is obtained in Prop. 3.1, in terms of the function,

$$H_\zeta(x, x') = \frac{d}{d\zeta} h_\zeta(x, x'), \quad x, x' \in \mathbb{X}. \quad (28)$$

The invariant pmf π_ζ for P_ζ is regarded as the equilibrium state for the mean-field model (1), with respect to the constant input value $\zeta_t \equiv \zeta$. The linearization about this equilibrium is described in Prop. 3.1. The proof is omitted since it is minor generalization of [4, Prop. 2.4].

Proposition 3.1. *The linearization of (1) at a particular value ζ is the state space model with transfer function,*

$$G_\zeta(z) = C[Iz - A]^{-1}B \quad (29)$$

in which $A = P_\zeta^T$, $C_i = \tilde{\mathcal{U}}_\zeta(x^i)$ for each i , and

$$B_i = \sum_x \pi_\zeta(x) \mathcal{E}_\zeta(x, x^i), \quad 1 \leq i \leq d \quad (30)$$

■

Another representation of B is obtained based on the product $P_\zeta^\nabla = P_\zeta^\dagger P_\zeta$, where P_ζ^\dagger denotes the adjoint of P_ζ .

Proposition 3.2. *The derivative of the transition matrix can be expressed in terms of the function (28):*

$$\frac{\mathcal{E}_\zeta(x, x')}{P_\zeta(x, x')} = H_\zeta(x, x') - P_\zeta H_\zeta(x) \quad (31)$$

where $P_\zeta H_\zeta(x) = \sum_{x'} P_\zeta(x, x') H_\zeta(x, x')$ for $x \in \mathbb{X}$. In the special case in which $H_\zeta(x, x')$ is independent of x , the entries of the vector B can be expressed,

$$B_i = \pi_\zeta(x^i) [H_\zeta(x^i) - P_\zeta^\nabla H_\zeta(x^i)] \quad (32)$$

Proof. For each x, x' ,

$$\frac{d}{d\zeta} \log(P_\zeta(x, x')) = H_\zeta(x, x') - \frac{d}{d\zeta} \Lambda_{h_\zeta}(x)$$

where we have used the definition $H_\zeta = \frac{d}{d\zeta} h_\zeta$. The derivative is computed using (5), giving

$$\begin{aligned} &= \frac{d}{d\zeta} \log\left(\sum_{x'} P_0(x, x') \exp(h_\zeta(x, x'))\right) \\ &= \exp(-\Lambda_{h_\zeta}(x)) \sum_{x'} P_0(x, x') \exp(h_\zeta(x, x')) H_\zeta(x, x') \\ &= \sum_{x'} P_\zeta(x, x') H_\zeta(x, x') \end{aligned}$$

which implies (31). If H_ζ depends only on x' then,

$$B_i = \sum_x \pi_\zeta(x) P_\zeta(x, x^i) [H_\zeta(x^i) - P_\zeta H_\zeta(x)]$$

We can write $\pi_\zeta(x) P_\zeta(x, x^i) = \pi_\zeta(x^i) P_\zeta^\top(x^i, x)$ to obtain

$$\begin{aligned} B_i &= \sum_x \pi_\zeta(x^i) P_\zeta^\top(x^i, x) [H_\zeta(x^i) - P_\zeta H_\zeta(x)] \\ &= \pi_\zeta(x^i) H_\zeta(x^i) - \pi_\zeta(x^i) \sum_x P_\zeta^\top(x^i, x) [P_\zeta H_\zeta(x)] \end{aligned}$$

□

3.2 Power spectral density and the positive real condition

In [10] the transfer function (3) was considered for a linearized mean-field model in continuous time. A representation of this transfer function used in this prior work admits a counterpart in the discrete-time setting.

The infinite series on the right hand side of (3) suggests that we require a probabilistic interpretation of the scalar $CA^k B$, where (A, B, C) are given in Prop. 3.1. This is achieved on defining $\tilde{U}_{B_\zeta}(x^i) = B_i / \pi_\zeta(x^i)$ for each i ; this is a function on \mathbf{X} whose mean is zero: $0 = \sum_x \pi_\zeta(x) \tilde{U}_{B_\zeta}(x)$.

Lemma 3.3. *Let \mathbf{X} denote a stationary realization of the Markov chain with transition matrix P_ζ , so that in particular, $X(k) \sim \pi_\zeta$ for each k . Then,*

$$CA^k B = \mathbb{E}[\tilde{U}_{B_\zeta}(X(0)) \tilde{U}_\zeta(X(k))] \quad (33)$$

Proof. We have by definition, $CA^k B = B^T P_\zeta^k C^T$, which can be expressed as the sum,

$$B^T P_\zeta^k C^T = \sum_{i,j} \pi_\zeta(x^i) \tilde{U}_{B_\zeta}(x^i) P_\zeta^k(x^i, x^j) \tilde{U}_\zeta(x^j)$$

This is equivalent to (33). □

With these identities in place we are ready to prove the passivity bound in Theorem 2.3.

Proof of Theorem 2.3 In the SPD solution without probabilistic constraints, it follows from the design rule (20) and the representation for the vector B in (32) that $\tilde{U}_{B_\zeta} = \tilde{U}_\zeta$. Lemma 3.3 gives the covariance interpretation,

$$CA^k B = \mathbb{E}[\tilde{U}(X(0)) \tilde{U}(X(k))]$$

Let $S^+(\theta)$ denote the power spectral density,

$$S^+(\theta) = \sigma_\zeta^2 + \sum_{k=1}^{\infty} \mathbb{E}[\tilde{U}(X(0)) \tilde{U}(X(k))] [e^{jk\theta} + e^{-jk\theta}]$$

where $\sigma_\zeta^2 = \mathbb{E}[\tilde{U}(X(0))^2]$. The bound thus follows from the definitions:

$$G_\zeta^+(e^{j\theta}) + G_\zeta^+(e^{-j\theta}) = S^+(\theta) + \sigma_\zeta^2 \geq \sigma_\zeta^2$$

■

4 Examples

4.1 Rational pools

The load in this case is a pool pump used to maintain water quality in a residential pool. The pump is assumed to consume 1 kW of power when operating. In the nominal model it is assumed that it runs for between 8 and 14 hours per day. In the original model of [4], the state space was taken to be the finite set,

$$\mathbf{X} = \{(m, k) : m \in \{\oplus, \ominus\}, k \in \{1, \dots, \mathcal{I}\}\} \quad (34)$$

where $\mathcal{I} > 1$ is an integer. For the i th load, if $X(t)^i = (\oplus, k)$, this means that the pool pump is on at time t , and has remained on for the past k time units.




Figure 4: State transition diagram for pool pump model.

In this paper we take the same state space, but with a different interpretation of each state: Here as in [5] we employ geometric sampling, so that the nominal state transition matrix is of the form (24). The state transition diagram for P_0 is shown in Fig:ppp.

In the experiments that follow, the transition matrix S_0 is the model with 12-hour cleaning cycle from [4], in which $\mathcal{I} = 48$, and hence $d = |\mathbf{X}| = 96$. It is assumed that the BA sends a signal every five minutes, and that the geometric sampling parameter is $\gamma = 1/6$. Consequently, the state of each load changes every 30 minutes on average.




Figure 5: $H_0 = \mathcal{H}(P_0)$ in three designs for the pool pump model.

There is no need to model uncertainty from nature. Hence the function $H = \mathcal{H}(P)$ will depend only on its second variable: $H(x, x') = H(x')$ for all (x, x') .

In each of these experiments the nominal transition matrix P_0 was taken as an input to the algorithm, and not S_0 . The resulting transition matrix P_ζ is of the form given in (25), in which the sampling rate is state-dependent for non-zero ζ . For the SPD solution it was necessary to use P_0 as the input: It can be shown that the transition matrix P_0^∇ is irreducible, but $S_0^\nabla := S_0^\nabla S_0$ is *not* irreducible in this example. Recall from Theorem 2.3 that irreducibility is required to ensure the existence of the SPD solution.

Fig. 5 shows the function $H_0 = \mathcal{H}(P_0)$ obtained in three sets of experiments. The normalization in the myopic design uses $H_0 = 30(\mathcal{U} - 1/2)$. This is equivalent to using \mathcal{U} since adding a constant

does not impact R_ζ , and the multiplication by 30 only scales ζ . A comparison of the three transfer functions G^+ obtained through a linearization at $\zeta = 0$ is shown in Fig. 6. The myopic design appears to be preferable to the two others: in particular, the phase plot stays nearest to zero in this design.




Figure 6: Bode plots obtained from three designs for the pool pump model.

One drawback with the myopic design is that we have no basis for analysis. Difficulties are also observed in numerical experiments if we allow $|\zeta|$ to take on values far from zero. Of the three designs, the IPD is found to be the most numerically stable in all of the experiments considered, in the sense that the dynamics change predictably with ζ , and the linearization about each value of ζ are nearly the same for a large range of ζ .

In the next subsection we move to a different class of loads in which we cannot ignore exogenous randomness. We consider in greater depth the difference between the myopic and IPD outcomes for a wider range of ζ . It is found again that the two designs are very similar locally (for $\zeta \sim 0$), but the myopic design is numerically unstable for ζ outside of a small neighborhood of the origin.

4.2 Thermostatically controlled loads

A thermostatically controlled load (TCL) such as a water heater, refrigerator, or air-conditioner is a device for which temperature control is achieved using a dead-band. It is assumed here that power consumption takes just two values (on or off). To simplify discussion, attention is directed to a cooling device, such as a residential refrigerator.

We begin with a noise free model, described as the controlled linear system,

$$\Theta(k+1) = \Theta(k) + (1 - \varrho)(\Theta_a - \Theta(k) - m(k)\Theta_g), \quad (35)$$

where $m(k) = 1$ or 0 indicates if the unit is on or off, Θ_a is ambient temperature, ϱ depends on the dynamics of the load, and Θ_g the depends on the physics of the device. One example considered in [15] is an air conditioner for which $\Theta_g = RP_{\text{trans}} = 2 \times 14 = 28$ (this and other parameters are summarized in Table 1).

A signal from the BA is broadcast at 20 second intervals. The parameter ϱ is obtained based on this sampling time, and the product of thermal resistance and capacitance: $\varrho = \exp[-20/(RC)]$, with RC also in units of seconds. The value $RC = 4$ obtained from Table 1 is in units of hours — on scaling to seconds we obtain,

$$\varrho = e^{-h/(RC)} = e^{-20/(4 \times 60^2)} = e^{-1/720} \approx 1 - 1/720$$

Parameter	Meaning	Value
Θ_{set}	set temperature set-point	20°C
$[\Theta_{\text{min}}, \Theta_{\text{max}}]$	temperature dead-band	$20 \pm 0.5^\circ\text{C}$
Θ_a	ambient temperature	32°C
R	thermal resistance	$2^\circ\text{C}/\text{kW}$
C	thermal capacitance	$2 \text{ kWh}/^\circ\text{C}$
P_{trans}	energy transfer rate	14 kW

Table 1: Homogeneous air conditioner parameters — mean data from Table 4.1 of [15].

This model is based on the physics of heating and cooling, but the dynamics are accurately captured by a constant drift model:

$$\Theta(k+1) = \Theta(k) - m(k)\delta_- + (1 - m(k))\delta_+ \quad (36)$$

With drift parameters δ_\pm carefully chosen, the behavior of the two models is barely distinguishable. The deterministic model (36) is the basis of a stochastic model,

$$\Theta(k+1) = \Theta(k) - m(k)\delta_- + (1 - m(k))\delta_+ + \Delta(k+1) \quad (37)$$

in which Δ is a zero-mean, i.i.d. sequence. In the experiments that follow the sampling time was taken to be 2 seconds, and Δ was taken to be Gaussian with variance 10^{-6} (the small variance is justified with this fast sampling rate).




Figure 7: State transition diagram for a TCL providing cooling.

A Markov chain model can be constructed with state $X(k) = (m(k), X_n(k))$ evolving on $\mathbf{X} = \mathbf{X}_u \times \mathbf{X}_n$, where $X_n(k) = \Theta(k)$ and $\mathbf{X}_u = \{\ominus, \oplus\}$; the interpretation is the same as in the pool filtration model, with the interpretation $m(k) = \oplus$ is the same as $m(k) = 1$ in (37). Temperatures are restricted to a lattice to obtain a finite state-space Markov chain. To obtain d states, assume that $d \geq 4$ is an even number, and discretize the interval $[\Theta_{\text{min}}, \Theta_{\text{max}}]$ into $d/2$ values as follows: $\mathbf{X}_n = \{\Theta_{\text{min}} + kT_\Delta : 0 \leq k \leq d/2 - 1\}$, in which the increments in the lattice are $T_\Delta = (\Theta_{\text{max}} - \Theta_{\text{min}})/(d/2 - 1)$.

A nominal randomized policy for \mathbf{m} defines the transition matrix R_0 . Following the notation of [4], the nominal transition matrix for \mathbf{m} is defined by

$$R_0(x, \oplus) = \begin{cases} 1 - p^\ominus(x_n) & x = (\oplus, x_n) \\ p^\oplus(x_n) & x = (\ominus, x_n) \end{cases} \quad (38)$$

$$R_0(x, \ominus) = \begin{cases} p^\ominus(x_n) & x = (\oplus, x_n) \\ 1 - p^\oplus(x_n) & x = (\ominus, x_n) \end{cases}$$

As in [4], the definition of p^\ominus is based on the specification of a cumulative distribution function F^\ominus defined on the interval $[\Theta_{\text{min}}, \Theta_{\text{max}}]$. This CDF is meant to model the statistics of the time

interval during which the unit is off, for the model with continuous state space. We define $p^\ominus(x_n) = 1$ for $x_n = \Theta_{\max}$, $p^\ominus(x_n) = F^\ominus(x_n)$ for $x_n = \Theta_{\min}$, and for all other values,

$$p^\ominus(x_n) = [F^\ominus(x_n) - F^\ominus(x_n - T_\Delta)]/[1 - F^\ominus(x_n - T_\Delta)]$$

In the experiments that follow, the general form taken for F^\ominus was chosen in the parameterized family,

$$F^\ominus(x_n) = \exp(-(\Theta_{\max} - x_n)^\rho/(2\sigma^\rho)), \quad \Theta_{\min} \leq x_n \leq \Theta_{\max},$$

with $\sigma, \rho > 0$. The values $T_\Delta = 0.05$, $\sigma = 0.02$ and $\rho = 0.75$ were used for p^\ominus and p^\oplus in the experiments surveyed here.




Figure 8: Trajectory of TCL model with geometric sampling $\gamma = 1/3$.

In addition, geometric sampling was applied: a family of models of the form (26) was constructed, in which R_0 was used to define S_0 in (24). The construction of a model of this form requires a different interpretation of the nature component of the state \mathbf{X}_n .

To obtain dynamics of the form (24), let $\{\tau_k\}$ denote the discrete renewal process in which $\tau_0 = 0$ and $\{\Delta_k = \tau_k - \tau_{k-1} : k \geq 1\}$ is i.i.d., with a geometric marginal:

$$\mathbb{P}\{\Delta_k > n\} = \gamma^n, \quad n \geq 0, k \geq 1.$$

The nature component of the state is constant on each discrete-time interval ($t : \tau_k \leq t < \tau_{k+1}$), with $X_n(t) = \Theta(\tau_k)$ on this interval. Given a nominal randomized policy for the input process $\mathbf{X}_u = \mathbf{m}$, the nominal transition matrix Q_0 can be estimated via Monte-Carlo based on a simulation of (37), or from measurements of an actual TCL.

A Markov model with this transition matrix would also require $m(t)$ constant on each of the intervals ($t : \tau_k \leq t < \tau_{k+1}$), $k \geq 0$. In simulations this constraint was violated occasionally since $m(t) = 1$ when $\Theta(t) > \Theta_{\max}$, and $m(t) = 0$ when $\Theta(t) < \Theta_{\min}$. This leads to modeling error that is small, provided γ is not too close to unity. Fig. 8 shows an example of the evolution of \mathbf{X}_n with $\gamma = 1/3$. The temperature never violates the dead-band constraint because of the constraints imposed on \mathbf{X}_u .

In this example the transition matrix $S_0^\nabla := S_0^\nabla S_0$ is irreducible, so that the SPD solution is computable. Because of exogenous randomness, there is no motivation for this approach: Theorem 2.3 guarantees a passive linearization only when $P_0 = R_0$. Moreover, numerical results using this method were not encouraging: The resulting family of transition matrices $\{P_\zeta\}$ is extremely sensitive to ζ .

The linearization about $\zeta = 0$ for the myopic design was similar to the IPD solution, but as seen in Figs. 9a and 9b, the behaviors quickly diverge for values beyond $|\zeta| = 3$.

In conclusion, although the transfer functions for the linearizations at $\zeta = 0$ are nearly identical, in the myopic design the input-output behavior is unpredictable for $|\zeta| > 3$. The input-output behavior for IPD is much closer to a linear system for a wider range of ζ . This is consistent with results from prior research [2, 4, 5].

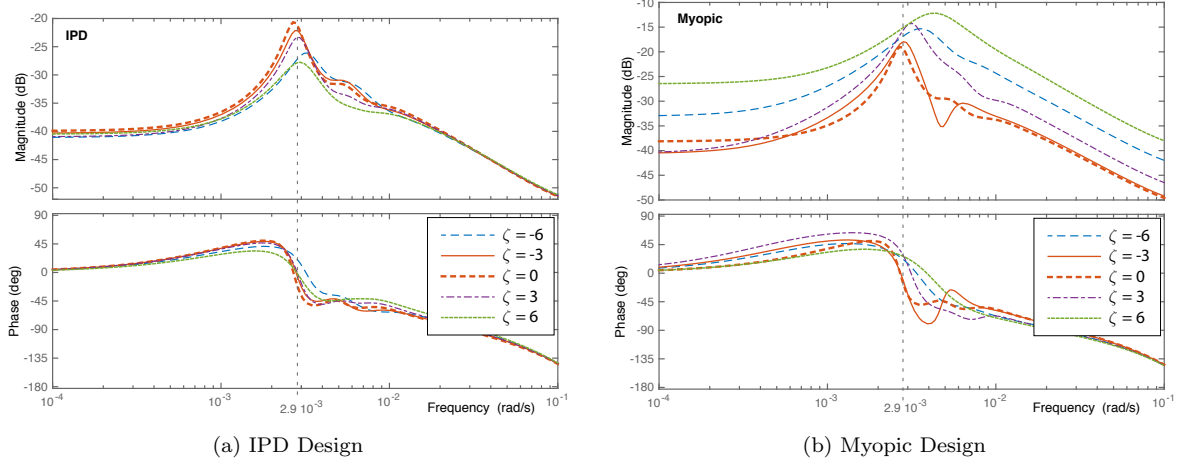


Figure 9: Bode plots for two designs, based on linearizations at five values of ζ . Bode plots for the myopic design are less reliable for $|\zeta| > 3$.

5 Conclusions

This paper has developed new approaches to distributed control for demand dispatch. There is much more work to do on algorithm design, and large-scale testing.

References

- [1] P. Barooah, A. Bušić, and S. Meyn, “Spectral decomposition of demand-side flexibility for reliable ancillary services in a smart grid,” in *Proc. 48th Annual Hawaii International Conference on System Sciences (HICSS)*, Kauai, Hawaii, 2015, pp. 2700–2709.
- [2] Y. Chen, A. Bušić, and S. Meyn, “Individual risk in mean-field control models for decentralized control, with application to automated demand response,” in *Proc. of the 53rd IEEE Conference on Decision and Control*, Dec. 2014, pp. 6425–6432.
- [3] H. Hao, Y. Lin, A. Kowli, P. Barooah, and S. Meyn, “Ancillary service to the grid through control of fans in commercial building HVAC systems,” *IEEE Trans. on Smart Grid*, vol. 5, no. 4, pp. 2066–2074, July 2014.
- [4] S. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, “Ancillary service to the grid using intelligent deferrable loads,” *IEEE Trans. Automat. Control*, vol. 60, no. 11, pp. 2847–2862, Nov 2015.
- [5] Y. Chen, A. Bušić, and S. Meyn, “State estimation for the individual and the population in mean field control with application to demand dispatch,” *CoRR and to appear, IEEE Transactions on Auto. Control*, 2015. [Online]. Available: <http://arxiv.org/abs/1504.00088v1>
- [6] S. Meyn, P. Barooah, A. Bušić, and J. Ehren, “Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in Florida,” in *Proceedings of the 52nd IEEE Conf. on Decision and Control*, Dec 2013, pp. 6946–6953.
- [7] J. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *IEEE Trans. Power Systems*, vol. 28, no. 1, pp. 430–440, 2013.

- [8] B. Sanandaji, H. Hao, and K. Poolla, “Fast regulation service provision via aggregation of thermostatically controlled loads,” in *47th Hawaii International Conference on System Sciences (HICSS)*, Jan 2014, pp. 2388–2397.
- [9] B. Biegel, L. Hansen, P. Andersen, and J. Stoustrup, “Primary control by ON/OFF demand-side devices,” *IEEE Trans. on Smart Grid*, vol. 4, no. 4, pp. 2061–2071, Dec 2013.
- [10] A. Bušić and S. Meyn, “Passive dynamics in mean field control,” in *Proc. 53rd IEEE Conference on Decision and Control*, Dec 2014, pp. 2716–2721.
- [11] E. Todorov, “Linearly-solvable Markov decision problems,” in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. Platt, and T. Hoffman, Eds. Cambridge, MA: MIT Press, 2007, pp. 1369–1376.
- [12] P. Guan, M. Raginsky, and R. Willett, “Online Markov decision processes with Kullback-Leibler control cost,” *IEEE Trans. Automat. Control*, vol. 59, no. 6, pp. 1423–1438, June 2014.
- [13] S. P. Meyn and R. L. Tweedie, *Markov chains and stochastic stability*, 2nd ed. Cambridge: Cambridge University Press, 2009, published in the Cambridge Mathematical Library. 1993 edition online.
- [14] A. Bušić and S. Meyn, “Respecting nature and nurture in Markov Decision Processes with Kullback–Leibler control cost,” 2015, In preparation.
- [15] J. Mathieu, “Modeling, analysis, and control of demand response resources,” Ph.D. dissertation, University of California at Berkeley, 2012.